


Impact on TRMM Products of Conversion to Linux

Erich Franz Stocker, NASA/GSFC Code 610.2
John Kwiatkowski, George Mason University



Reasons for Conversion


- **Since beginning of mission TRMM data products have been produced on an SGI**
 - 32 bit architecture
 - SGI Fortran and C compilers
 - HDF4 data format
- **Only have a single new SGI left**
- **SGI has indicated that would no longer support the older architectures and maintenance costs are high**
- **Due to budget reductions necessary to drop the support of the TSDIS software and transition to the Precipitation Processing System supporting GPM**
- **Preparation of TRMM V7 reprocessing and GPM V0 algorithms**



Why Worry




- **TRMM algorithm code contains many floating point calculations**
- **Floating point calculations are not always portable**
- **Floating point calculations are susceptible to**
 - Different hardware
 - Different operating systems
 - Different word size (e.g. 64 vs 32 bits)
- **In Fortran programs floating point calculations are even influenced by the compiler being used.**
- **Floating point round-off is affected by all the above**



Why PPS --




- **Hybrid not pure system**
 - PR L1A,L1B, L1C all produced on the current SGI
 - TRMM Realtime producing is on the SGI
 - PR L1 and Realtime products are big-endian while all others are little-endian
- **Still uses the TSDIS toolkit**
- **PPS database structure had to be augmented to deal with the TRMM special casing**
- **Produces TRMM V6 products**



Data Production Hardware




- **Beowulf cluster environment**
 - Penguin hardware
 - 64 bit Scyld Beowulf operating system (based on CentOS Linux)
- **Production**
 - 2 Dual core host nodes, AMD opteron, 8GB memory each (1 host active and 1 cold backup)
 - Run the compute nodes and handle outside network connection
 - 96 compute nodes
 - Single dual core AMD opteron each 6 GB memory
 - No disks (boot from host)
 - Panasas cluster file fast storage (processing buffer): 40TB
 - Storage manager server
 - Mini-cluster
 - Host node with two dual core AMD opteron 8 GB memory
 - 8 compute nodes with extra 1 GB network card
 - Manage 800 TB of RAID 6 SATA. Permanent archive storage
 - gcc C compiler and Intel Fortran 77/90 compiler



Integration and Test Hardware



- **Purpose**
 - Used for testing and integrating new or updated science algorithm code
 - Debugging
 - Initial multi-month test of new or updated algorithms
- **Beowulf clustered hardware**
- **Set up similar to the production cluster so a valid testing platform**
- **Hardware**
 - 1 dual core host node with two dual core AMD opterons, 8GB of memory
 - 32 compute nodes
 - Single core AMD opteron, 6 GB of memory
 - Diskless, booting from host
 - Panasas clustered file fast storage (processing buffer): 20 TB



Floating Point Issues



- **Floating point calculations**
- TRMM algorithms perform floating point calculations that have portability issues for porting from SGI to Linux Cluster. Rounding is generally machine dependent and/or compiler dependent as in the case of FORTRAN. This has implications for any decisions based on floating point values (e.g. thresholding, gridbox placement). This then impacts the final values of retrieved quantities, including rainfall rate.
 - Gridding with floating point latitude and longitude.
 - Scaling of values written into HDF products.
 - Algorithm decision trees:
 - Rain/no-rain detection.
 - Surface type detection (land/ocean/coast).



Geo-location Issues




The IFOVs of each instrument are geolocated onto the Earth ellipsoid at Level-1.

Latitude and longitude are stored as single precision floating point values in the HDF products.

While the changes are small and at the limit of single precision floating point accuracy, these differences do sometimes have an impact on the final retrievals.

Orbital maximum difference Latitude: $\pm 3.0 \times 10^{-6}$ degrees

Orbital maximum difference Longitude: $\pm 1.5 \times 10^{-5}$ degrees



2A21 Comparisons

The 2A21 product produces surface cross-sections and path integrated attenuation which are used in other PR and combined products. As a result comparison between an orbit produced on the SGI and Linux Cluster is instructive of the differences overall.

2A21.080101.57701.6L.HDF - /PANFS/data/sgi/2008/01/01/2A21.080101.57701.6.HDF

Name	MaxDiff	Frac MaxDiff	Data1 at MaxDiff	Data2 at MaxDiff	Frac MeanDif	#NonZero	Frac #NonZero
geolocation	0.00	0.	-34.0	-34.0	0.00e+00	0	0.e+00
sigmaZero	1.00	0.0009	1.15e+03	1.15e+03	1.38e-06	138	3.e-04
pathAtten	1.00	0.02	66.0	67.0	6.57e-08	7	2.e-05
reliabFlag	0.00	0.	1.99e+04	1.99e+04	0.00e+00	0	0.e+00
reliabFactor	0.000271	6.e-06	46.0	46.0	2.18e-10	17398	4.e-02
incAngle	0.00	0.	-182.	-182.	0.00e+00	0	0.e+00
rainFlag	0.00	0.	0.00	0.00	0.00e+00	0	0.e+00
TOTAL						17543	

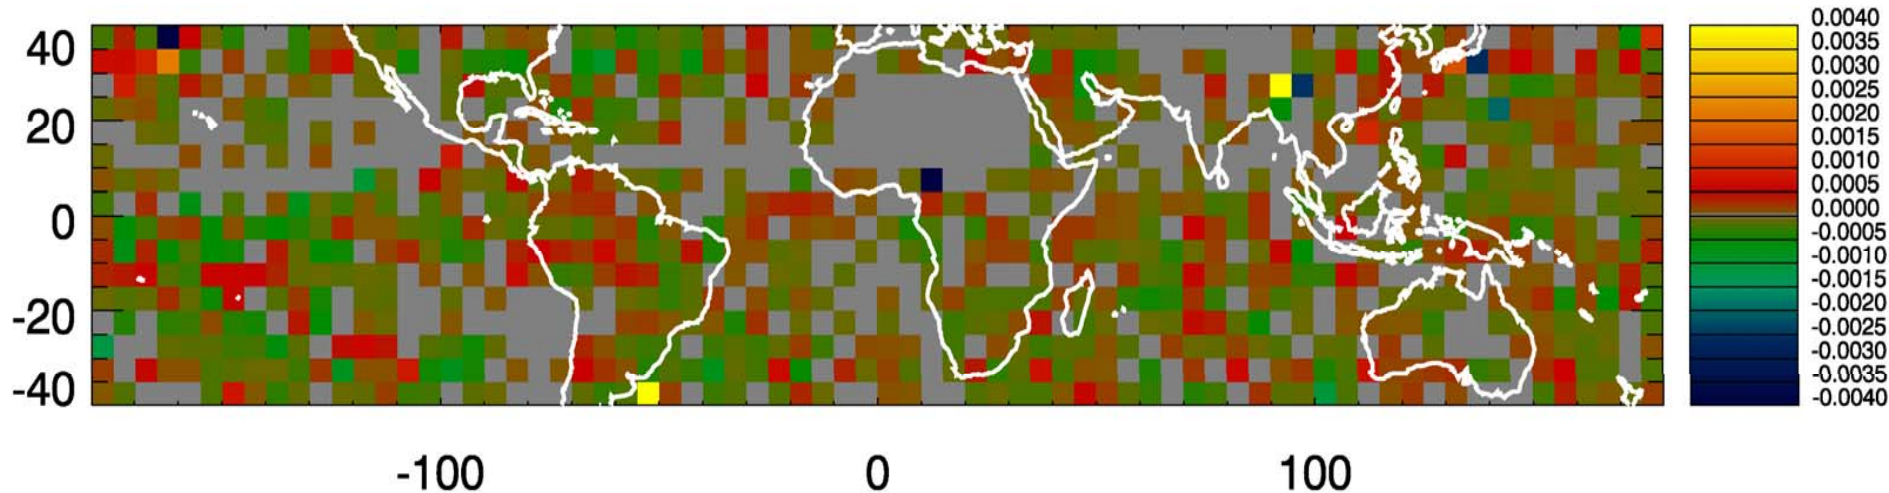
Surface cross section and path integrated attenuation are scaled to integers in HDF. A difference of 1.0 indicates a rounding of the last significant digit for those parameters.




Level 3 Comparisons



Level three products are good indicators of calculation differences between the two architectures. Round off differences can lead to a pixel being placed in a different grid box. The accumulation of level 2 values that are different also shows in the level 3. The following is a plot of differences for the PR level 3A25.



Maps of differences or % differences $(TSDIS-PPS)/TSDIS*100.0$ for various monthly, gridded averages. Grey indicates a value close to zero in the middle of the color scale, while **black** indicates no data (no rain). Positive differences are **red/yellow** while negative differences are **green/blue**.



Conclusions



- There are round-off differences in a few parameters in the same products produced on the different architectures
- **No scientifically significant differences exist between products produced on the different architectures**
- Had production started on the Linux cluster and moved to IRIX SGI, we would be using Linux as a reference and worried about the differences in the other direction
- If the system stands up to its 2nd Operational Acceptance test, we believe there is NO reason not to use it



Impact on TRMM data products of the conversion to Linux processing

E. Stocker (1), J. Kwiatkowski (2)

(1) NASA/GSFC Code 610.2 Greenbelt MD., (2) George Mason University, Fairfax Va.

In June 2008, TRMM data processing will be assumed by the Precipitation Processing System (PPS). This change will also mean a change in the hardware production environment from an SGI 32 bit IRIX processing environment to a Linux (Beowulf) 64 bit processing environment. This change of platform and operating system addressing (32 to 64) has some influence on data values in the TRMM data products. This paper will describe the transition architecture and scheduling. It will also provide an analysis of what the nature of the product differences will be. It will demonstrate that the differences are not scientifically significant and are generally not visible. However, they are not always identical with those which the SGI would produce.